DynaGuide: Steering Diffusion Polices with Active Dynamic Guidance

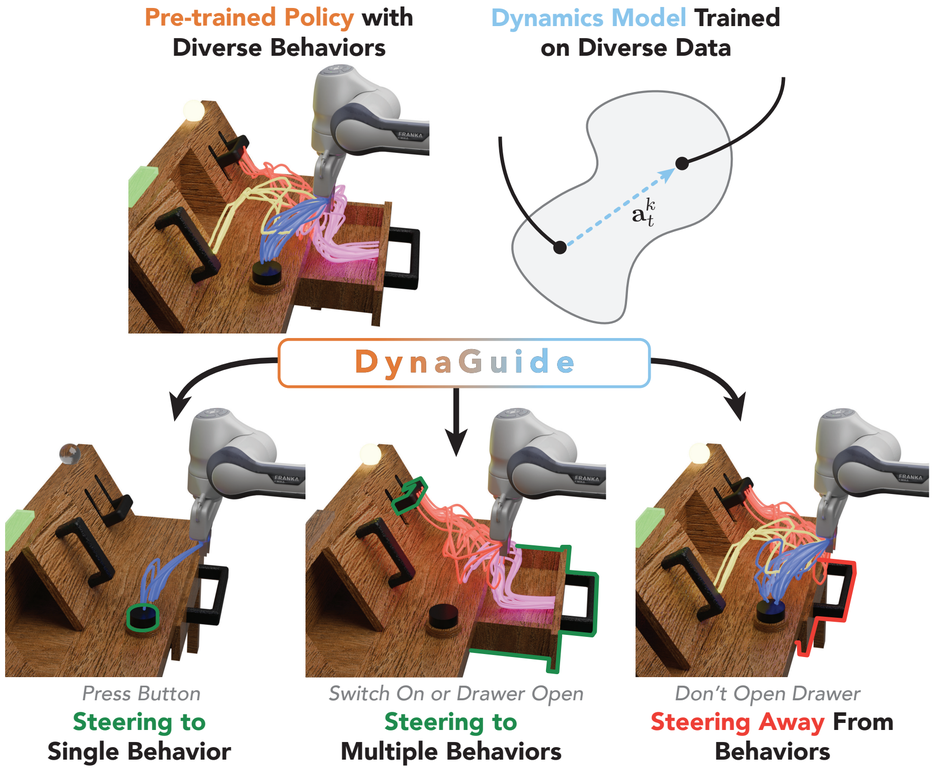

Maximilian Du, Shuran Song

Stanford Artificial Intelligence

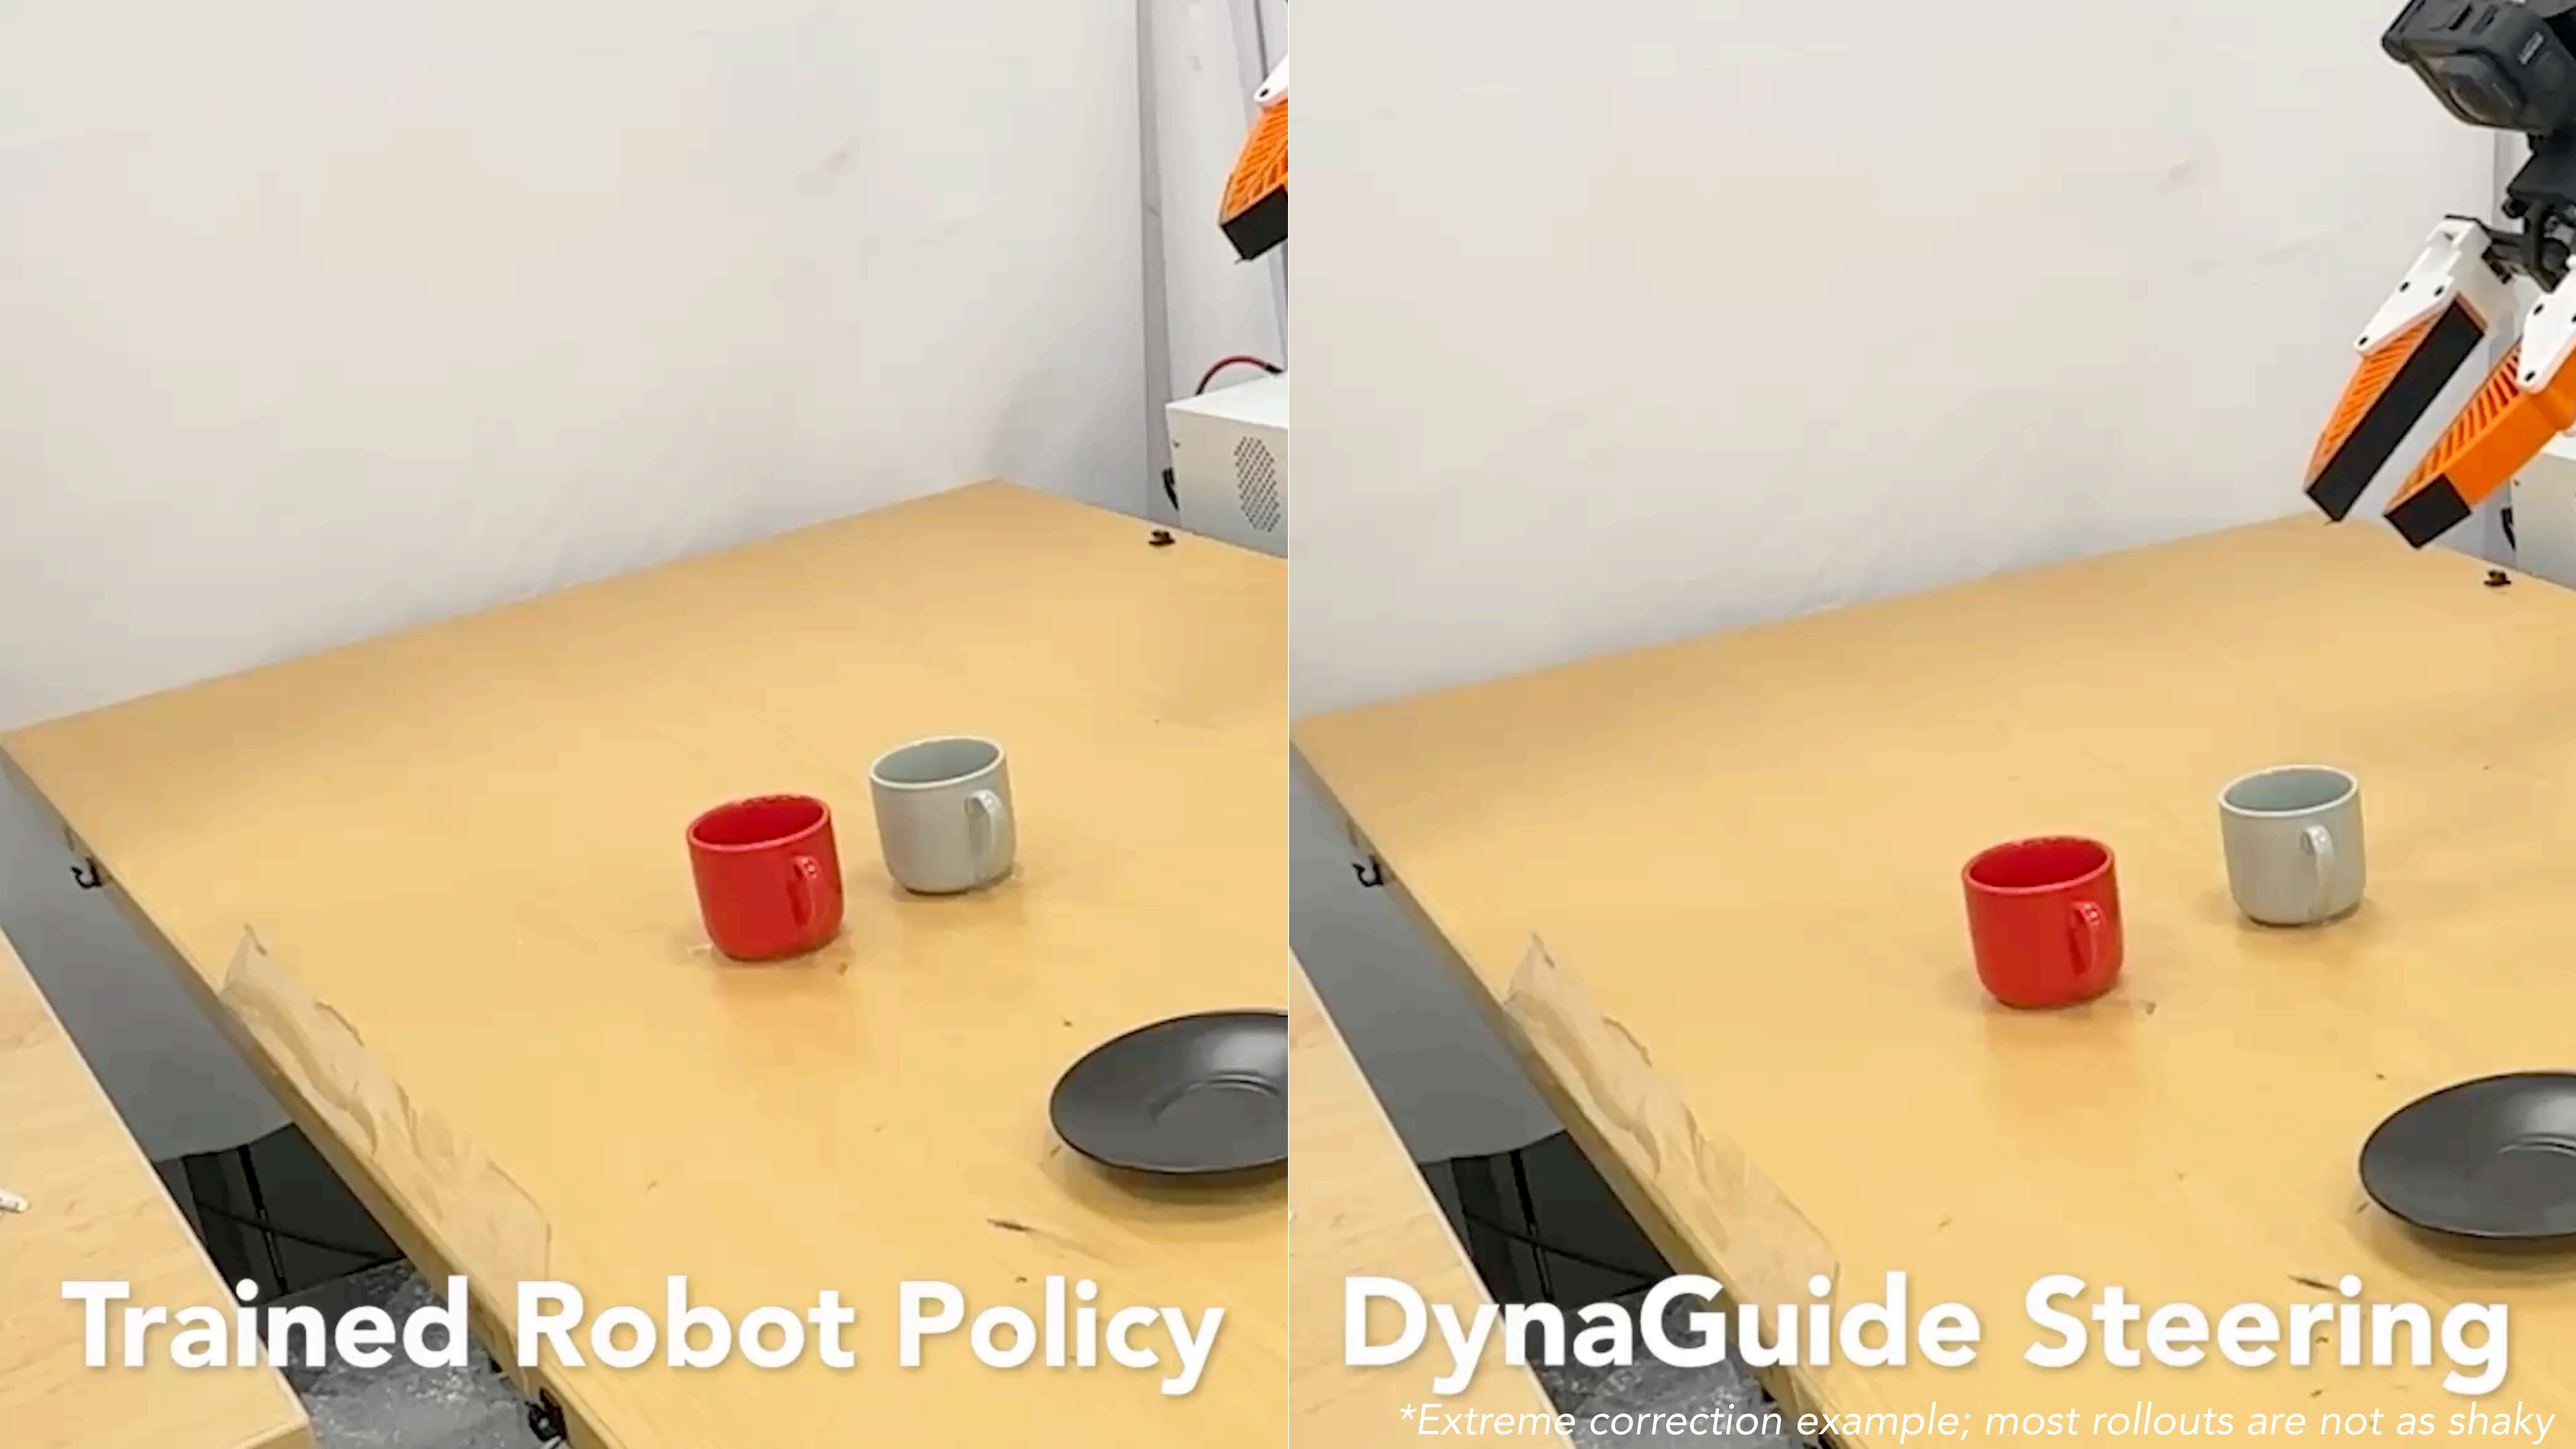
Trained Robot Policy
DynaGuide Steering
*Extreme correction example; most rollouts are not as shaky

How do we **steer** a robot policy?

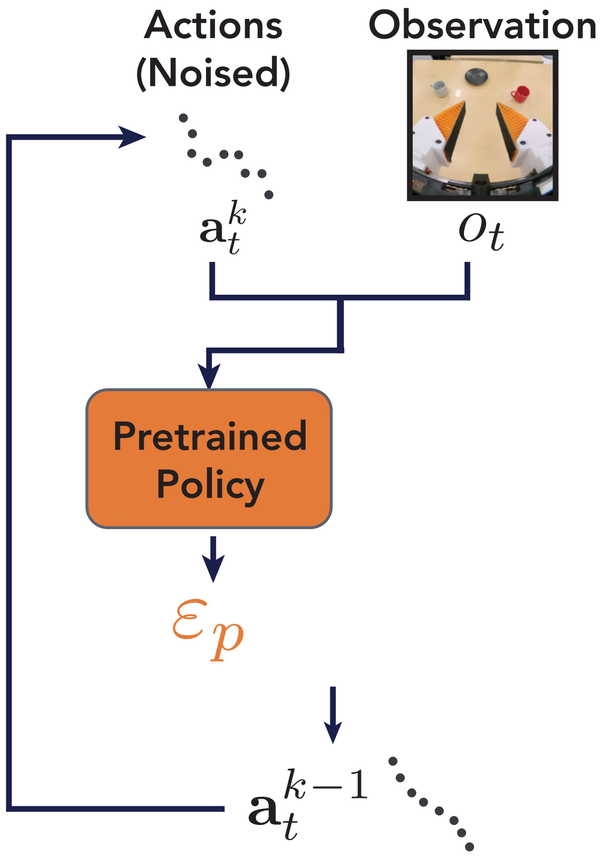

DynaGuide adds a **guidance signal** to the action denoising process in the diffusion policy

Paper Section: 3

How do we **steer** a robot policy?

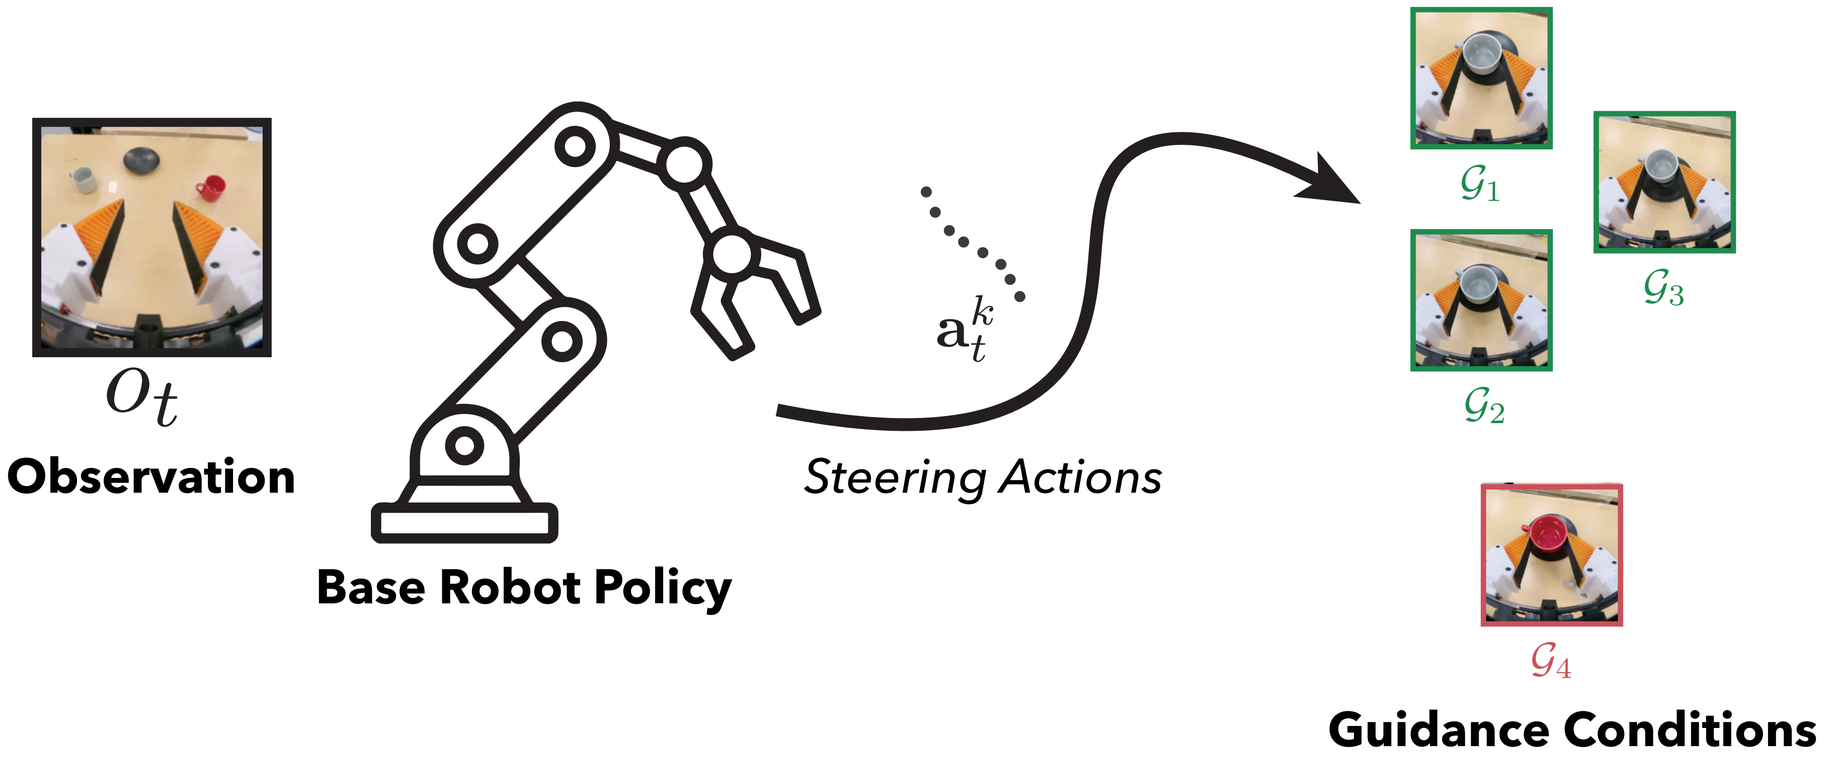

We need to influence a robot policy's **actions** to seek desired outcomes (or avoid undesired outcomes)

Paper Section: 3

How do we **steer** a robot policy?

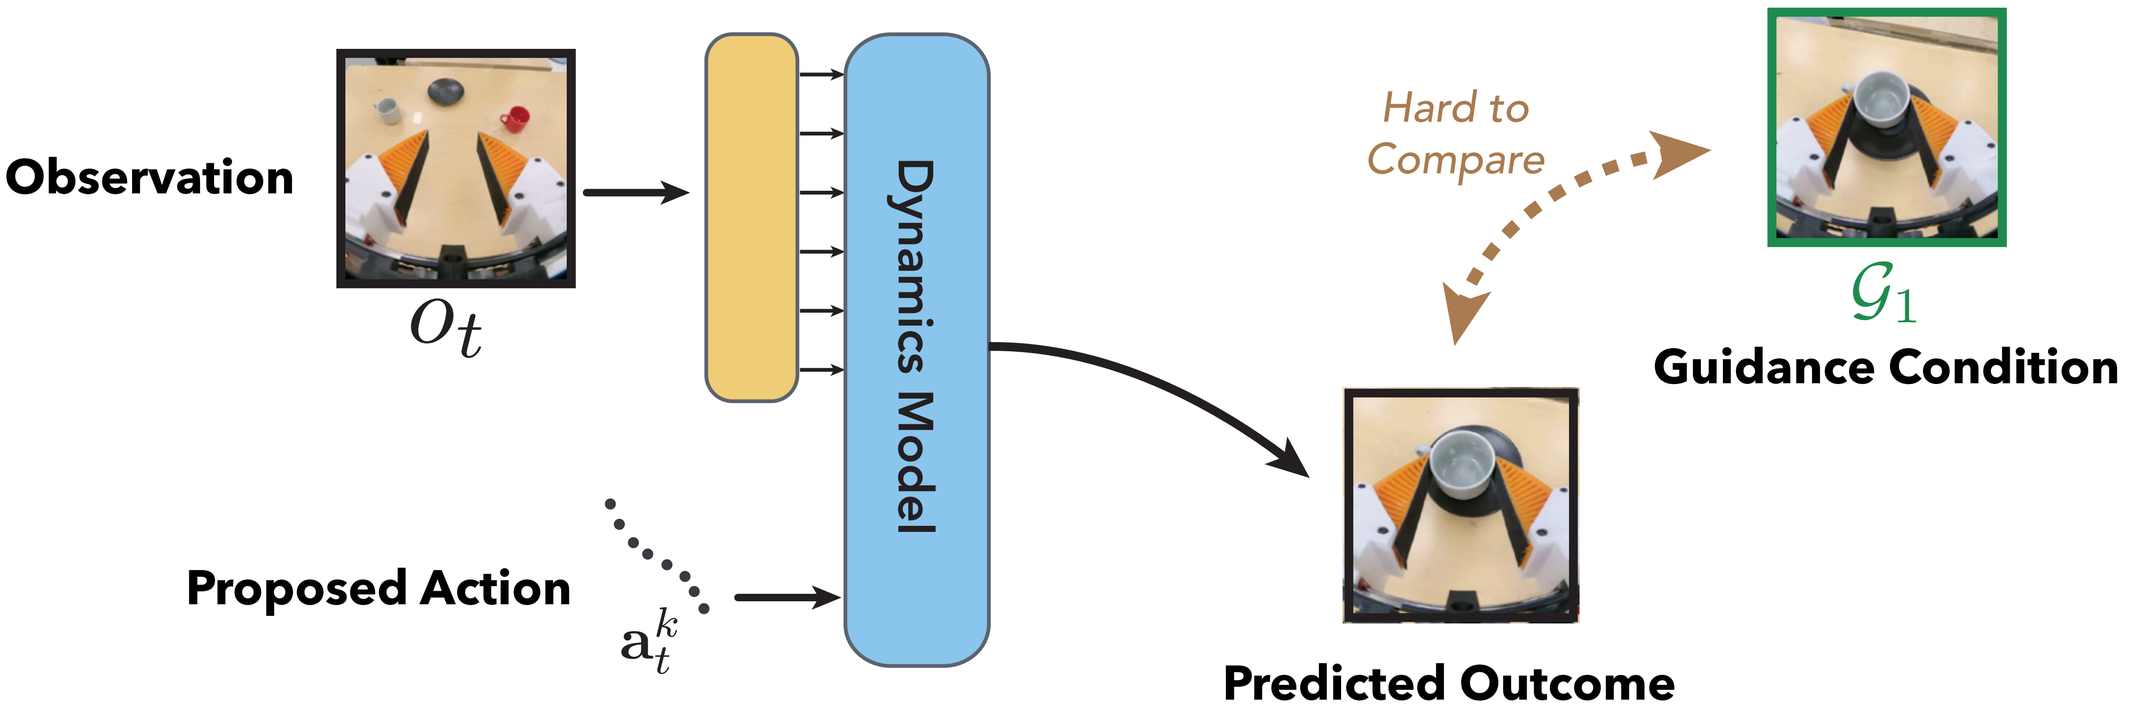

Steering towards outcomes means needing to predict the future. We do this with a **dynamics model.**

Paper Section: 3

How do we **steer** a robot policy?

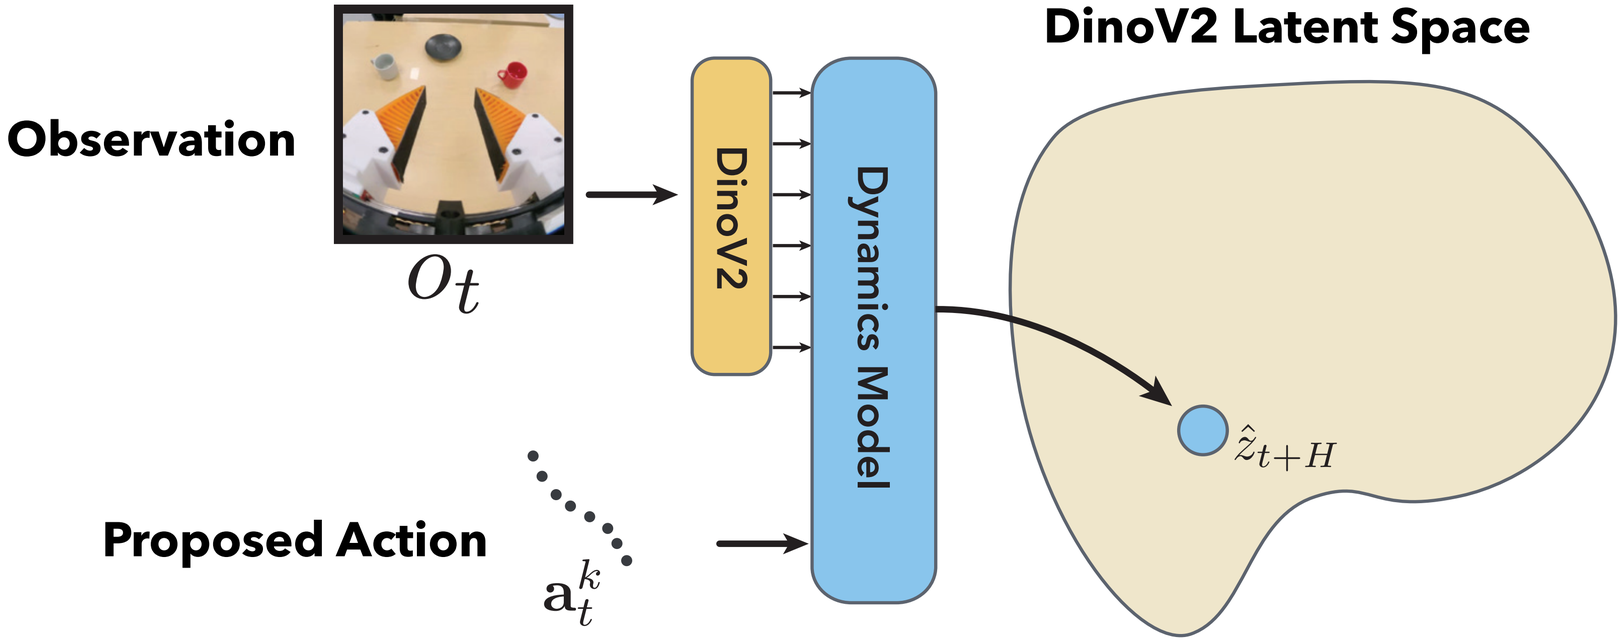

To work with visual observations / objectives, we train the dynamics model on a **latent representation.**

Paper Section: 3

How do we **steer** a robot policy?

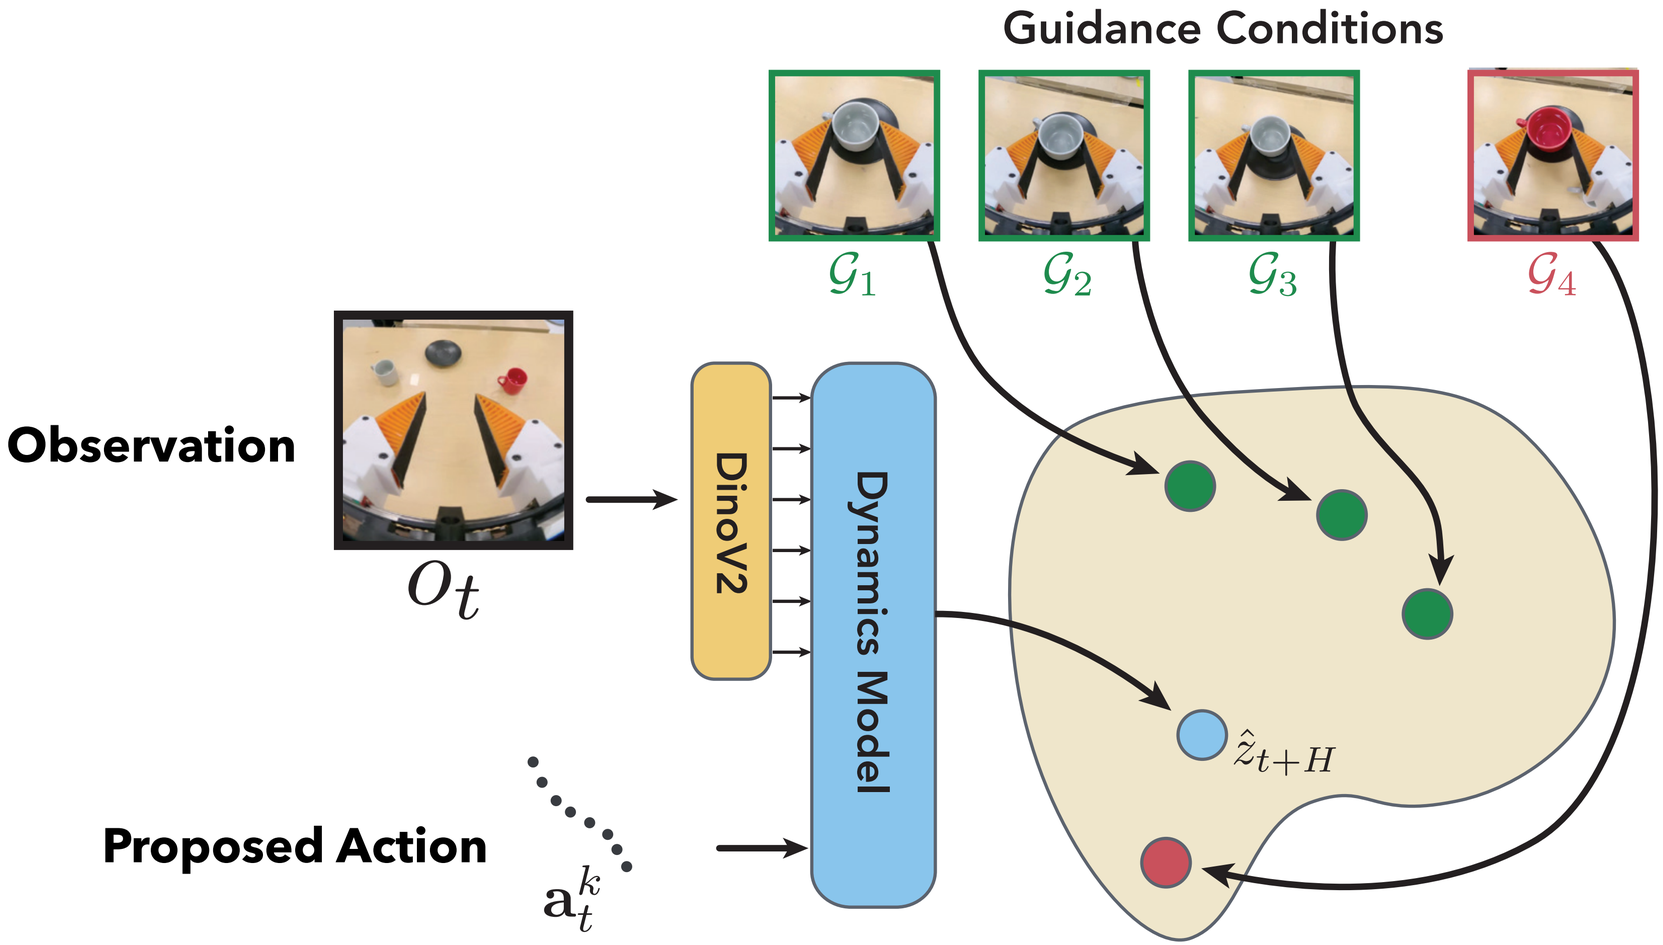

We can project desired / undesired outcomes into the same space, making comparisons straightforward (L2 distance)

Paper Section: 3

How do we **steer** a robot policy?

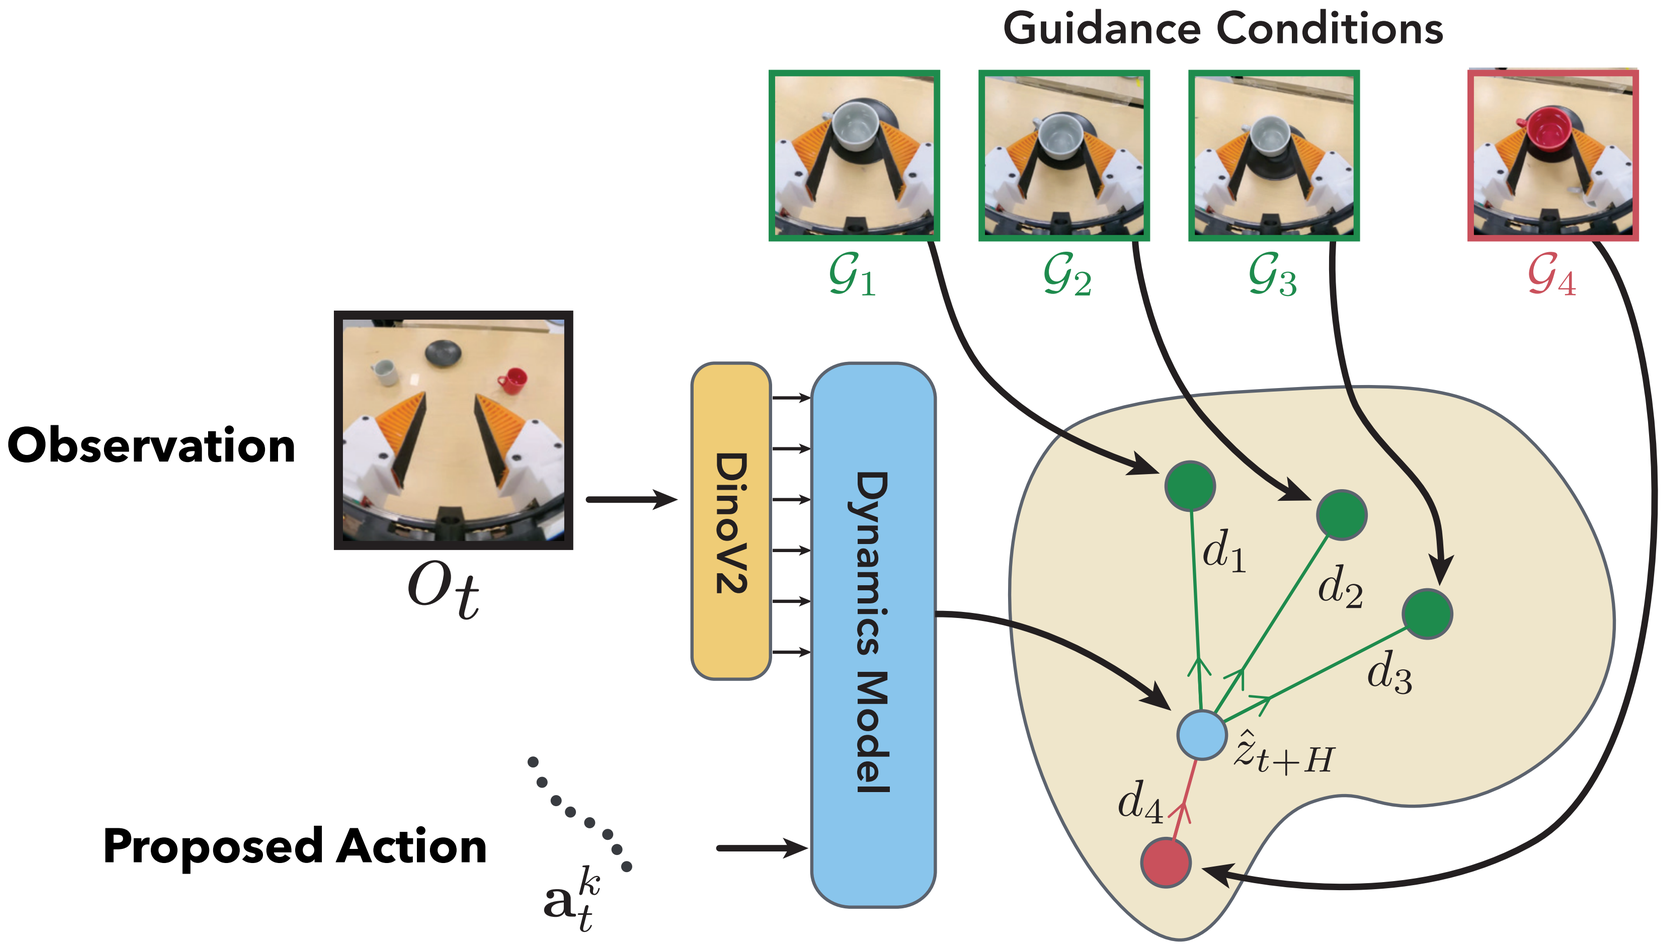

We can project desired / undesired outcomes into the same space, making comparisons straightforward (L2 distance)

Paper Section: 3

How do we **steer** a robot policy?

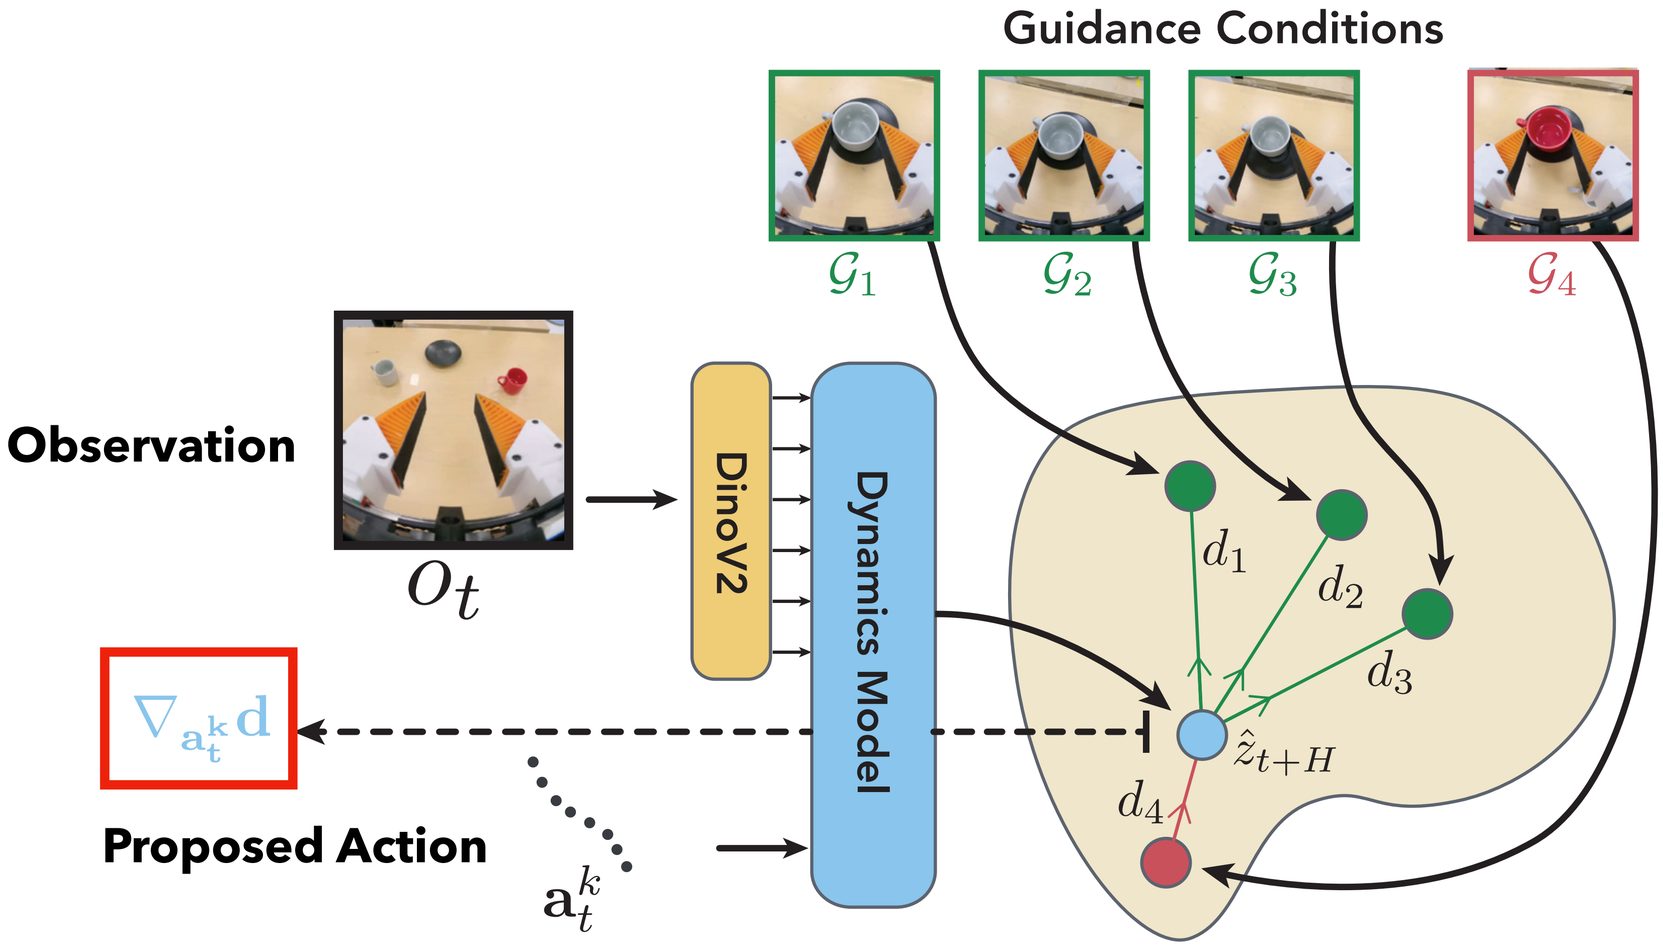

Backpropagate the distances through the dynamics model to get the action gradient

Paper Section: 3

How do we **steer** a robot policy?

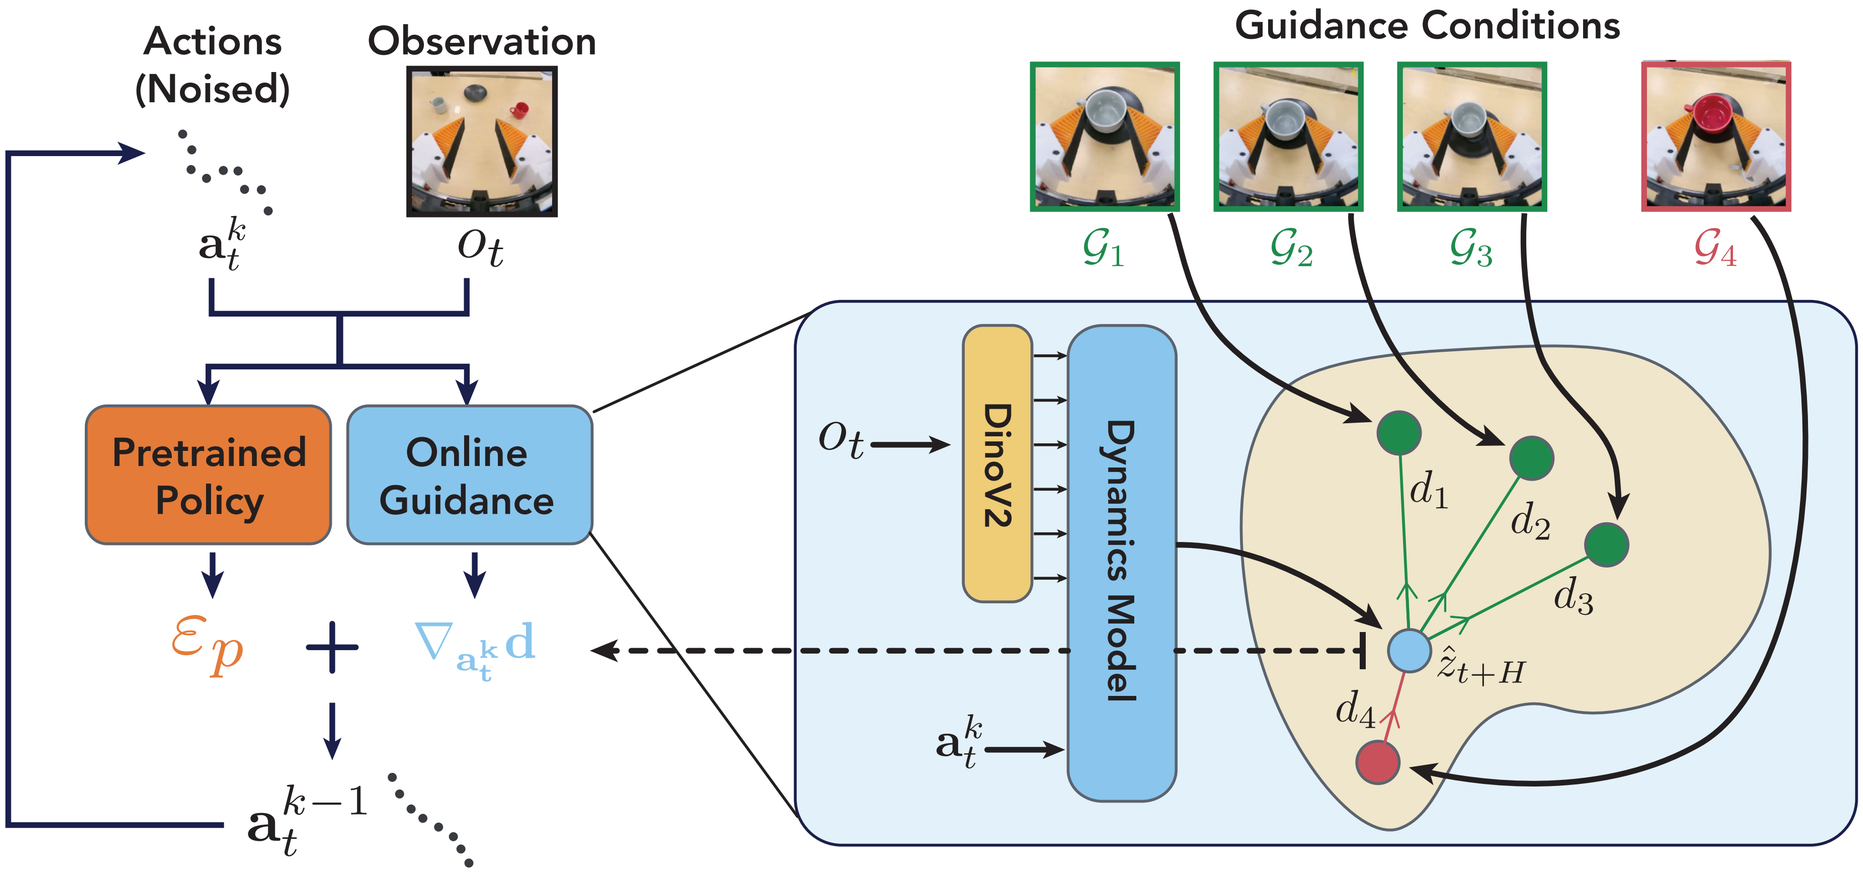

Combine the **guidance gradient** with the **policy noise estimation** to guide the action generation

Paper Section: 3

How do we **steer** a robot policy?

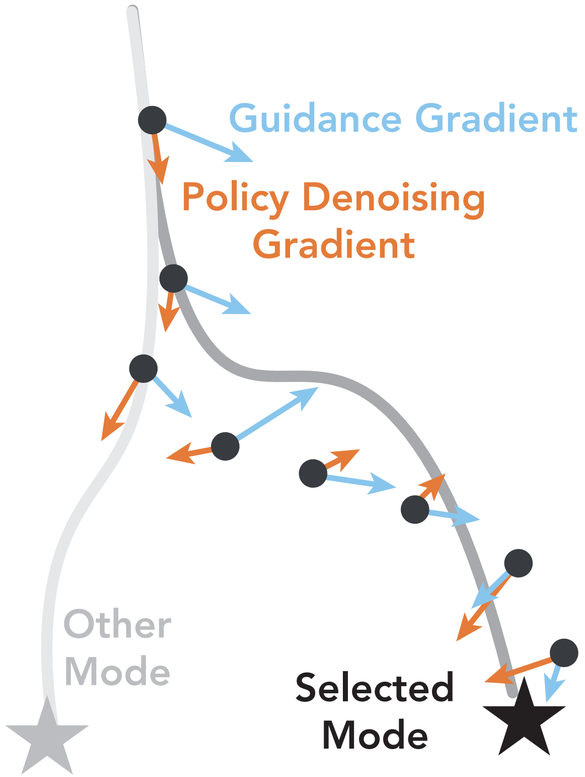

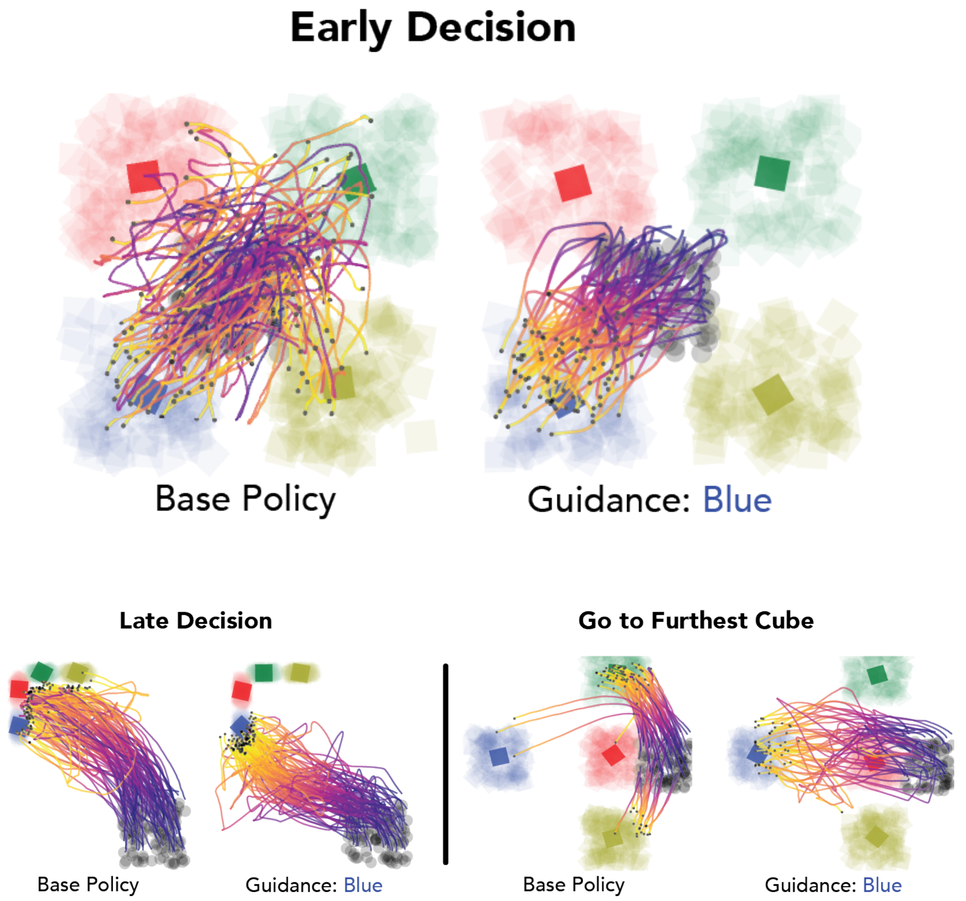

Combined influence will find the modes of the base policy that best satisfy the guidance conditions

Paper Section: 3, A.1

DynaGuide: CALVIN Simulated Experiments

1) Does DynaGuide successfully steer behavior?
2) How robust is DynaGuide to lower quality guidance conditions?
3) Can DynaGuide steer towards complex objectives?
4) Can DynaGuide enhance underrepresented behaviors in the base policy?

CALVIN: Experimental Setup

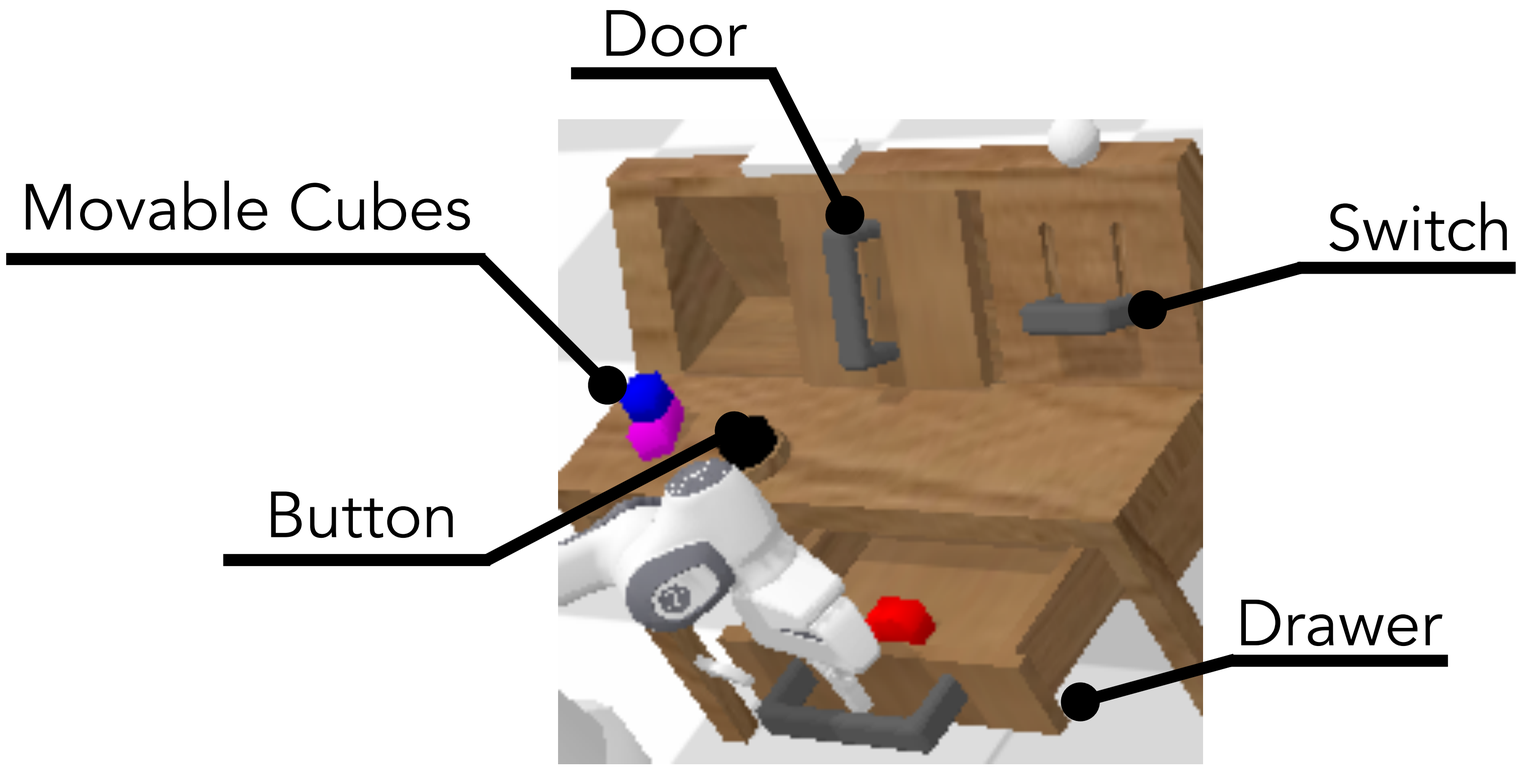

Evaluate DynaGuide and baselines on ability to steer behaviors in the CALVIN simulation, using human-collected play data provided by CALVIN.

Paper Section: 4

1) Does DynaGuide successfully steer behavior?

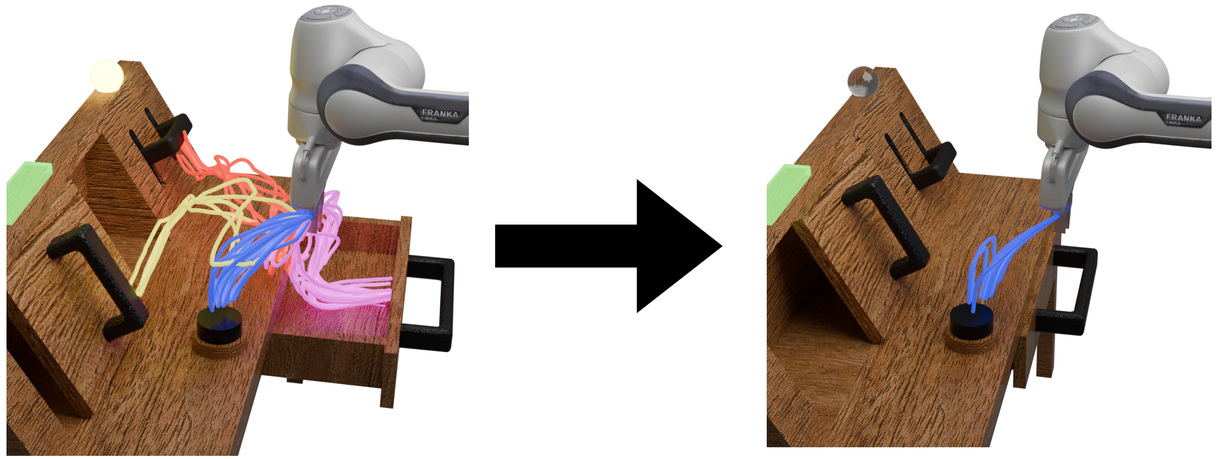

Diverse Base Policy **Specific Object Interaction**

Guidance Condition Examples (For switch on)

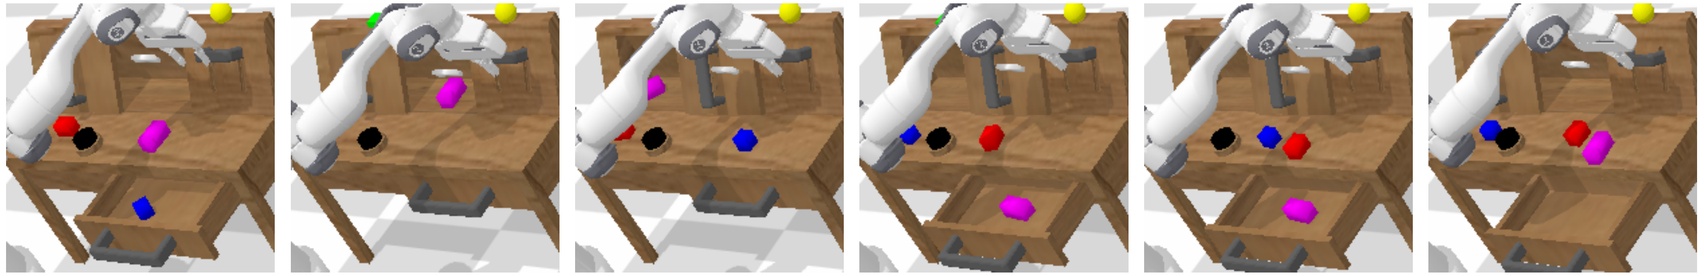

Investigate DynaGuide's ability to select one behavior from a diverse base policy

1) Does DynaGuide successfully steer behavior?

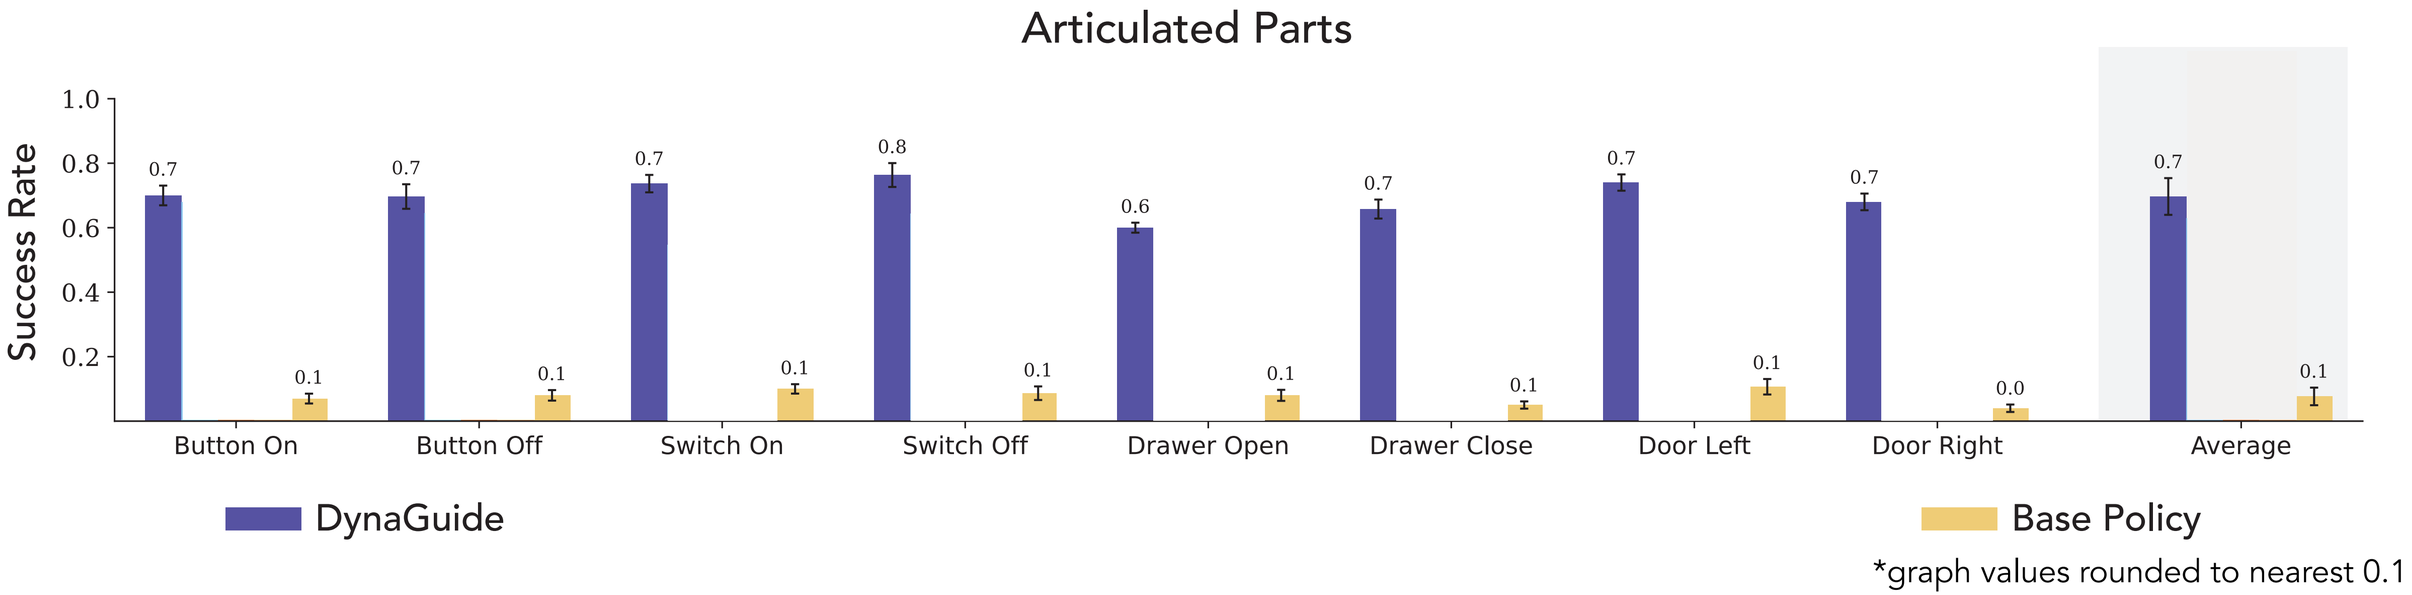

*graph values rounded to nearest 0.1

DynaGuide is ~70% successful on average at steering the **base policy**

Paper Section: 4.1

2) How robust is DynaGuide to lower quality guidance conditions?

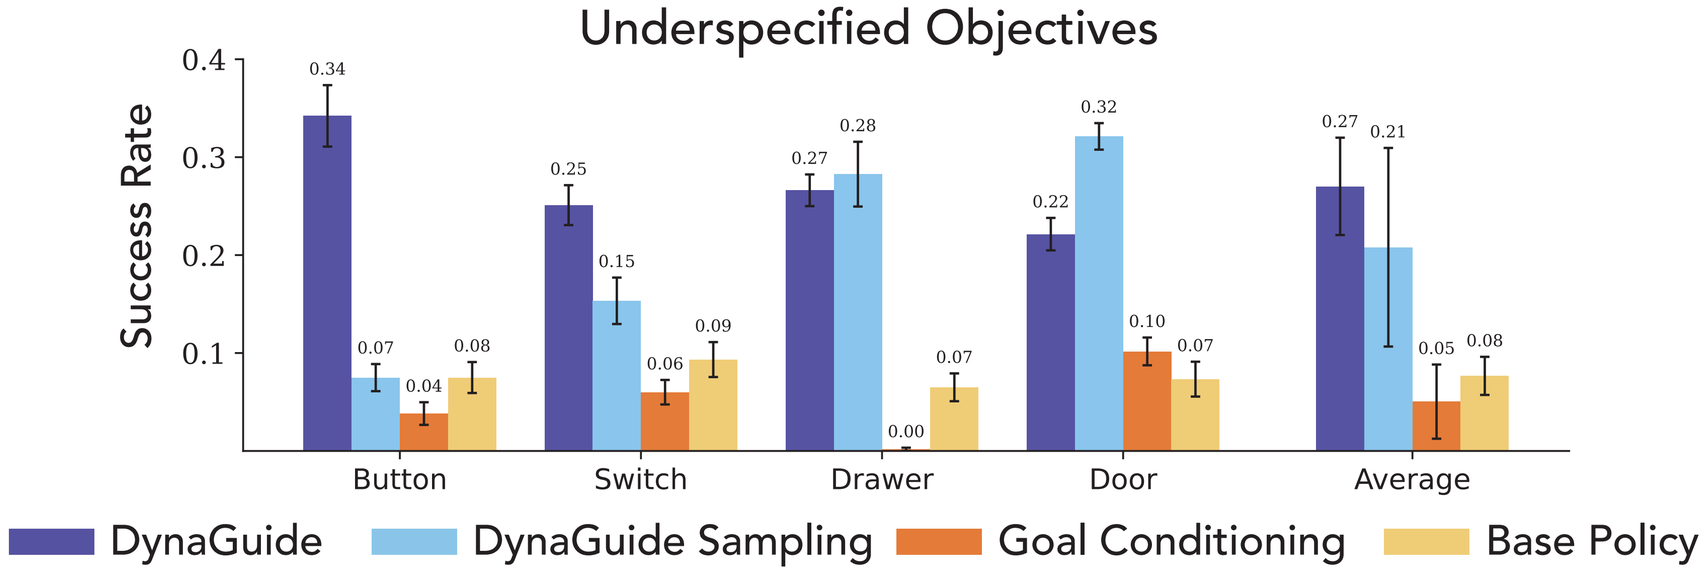

Guidance Conditions with Robot Randomization (Open Drawer)

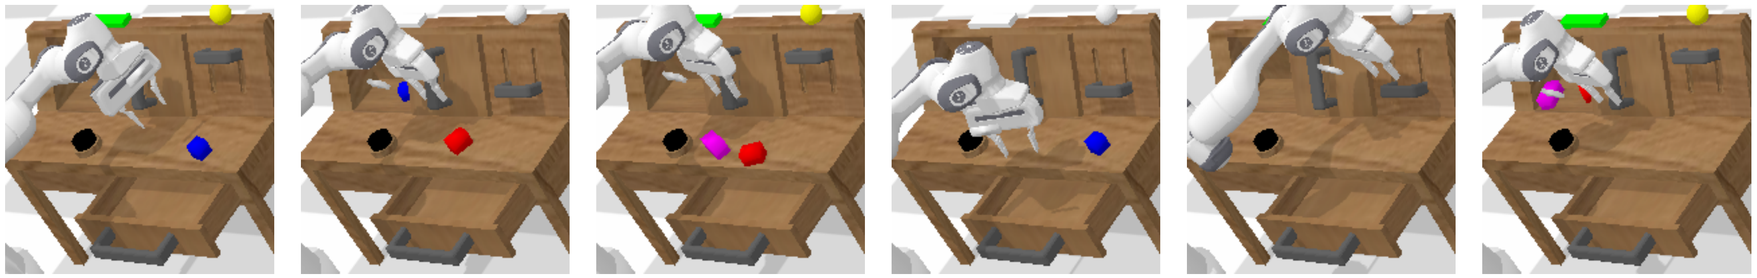

DynaGuide outperforms **goal conditioning** in robustness to noisy guidance conditions

Paper Section: 4.2

3) Can DynaGuide steer towards complex objectives?

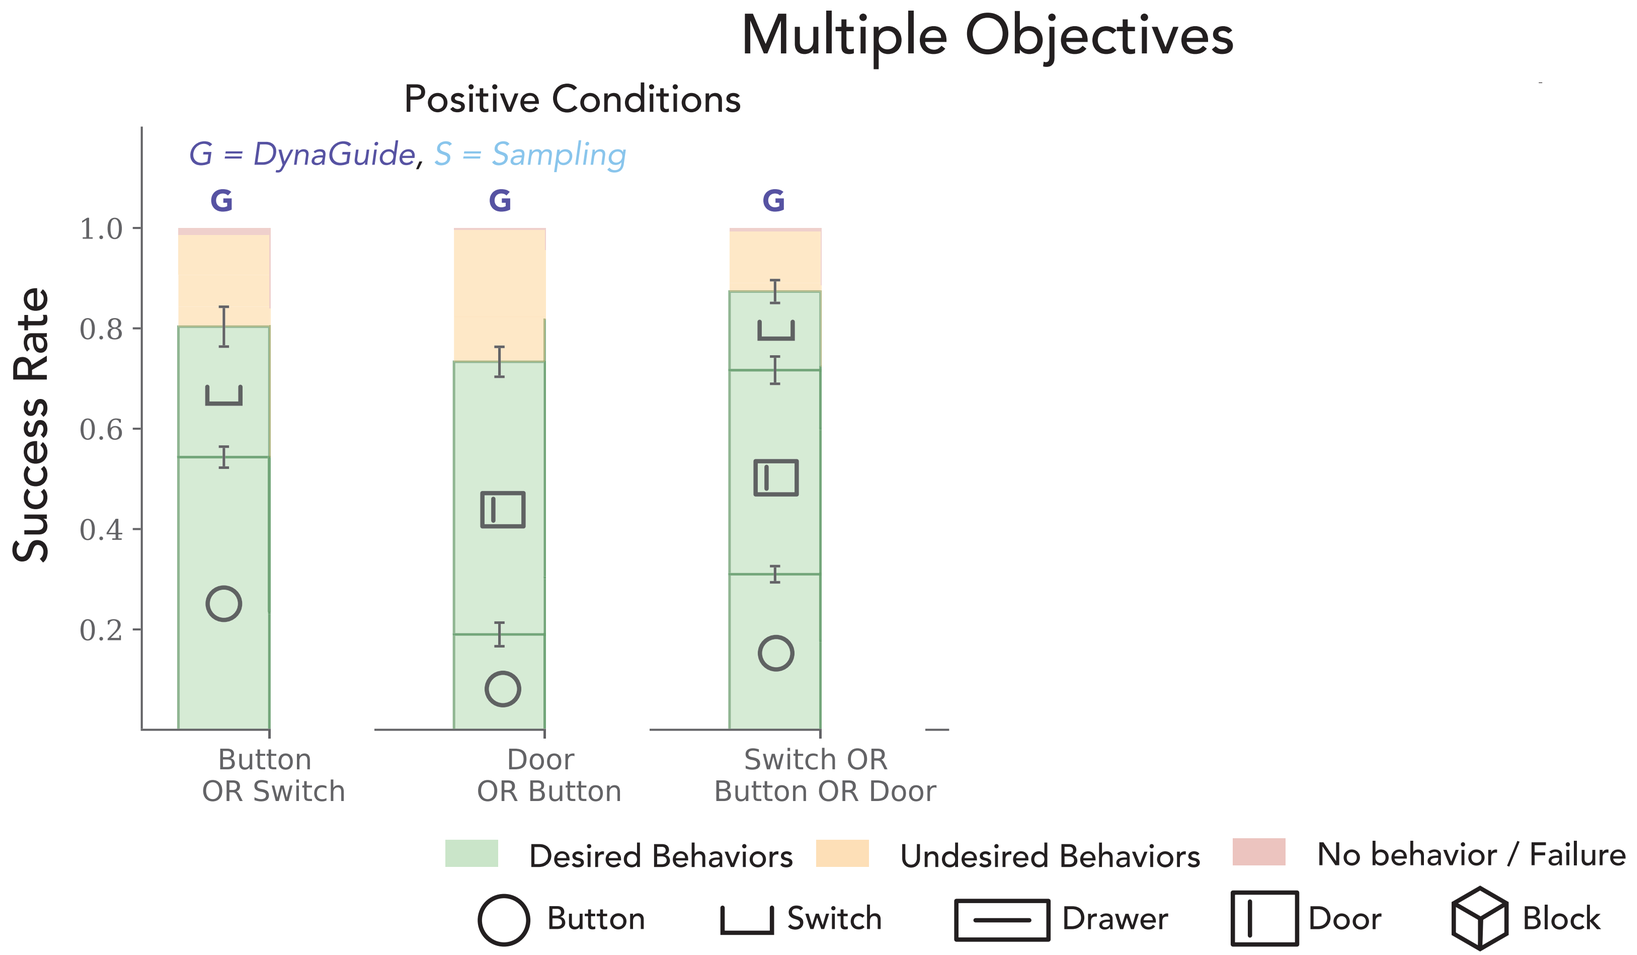

DynaGuide can steer policies towards multiple behaviors and avoid behaviors.

Paper Section: 4.3

4) Can DynaGuide enhance underrepresented behaviors in the base policy?

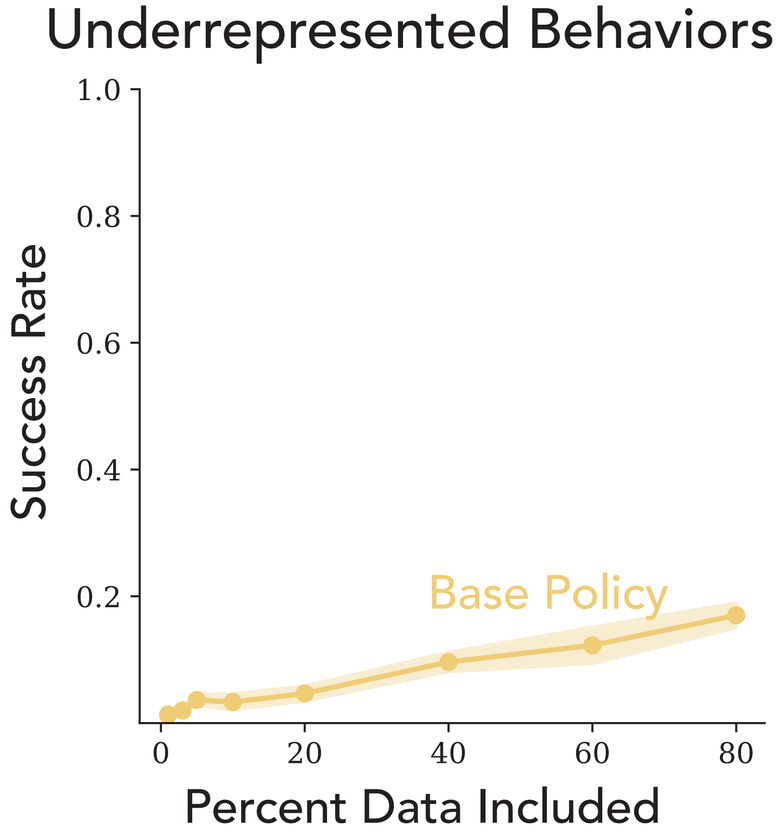

DynaGuide's active guidance enhances rare behaviors in the **base policy**, more than **sampling** approaches

Paper Section: 4.4

DynaGuide: Real Robot Experiments

1) Can DynaGuide steer an off-the-shelf real robot policy?
2) Can DynaGuide seek underrepresented behaviors on a real robot?

Paper Section: 4.5

DynaGuide: Real Robot Experiments

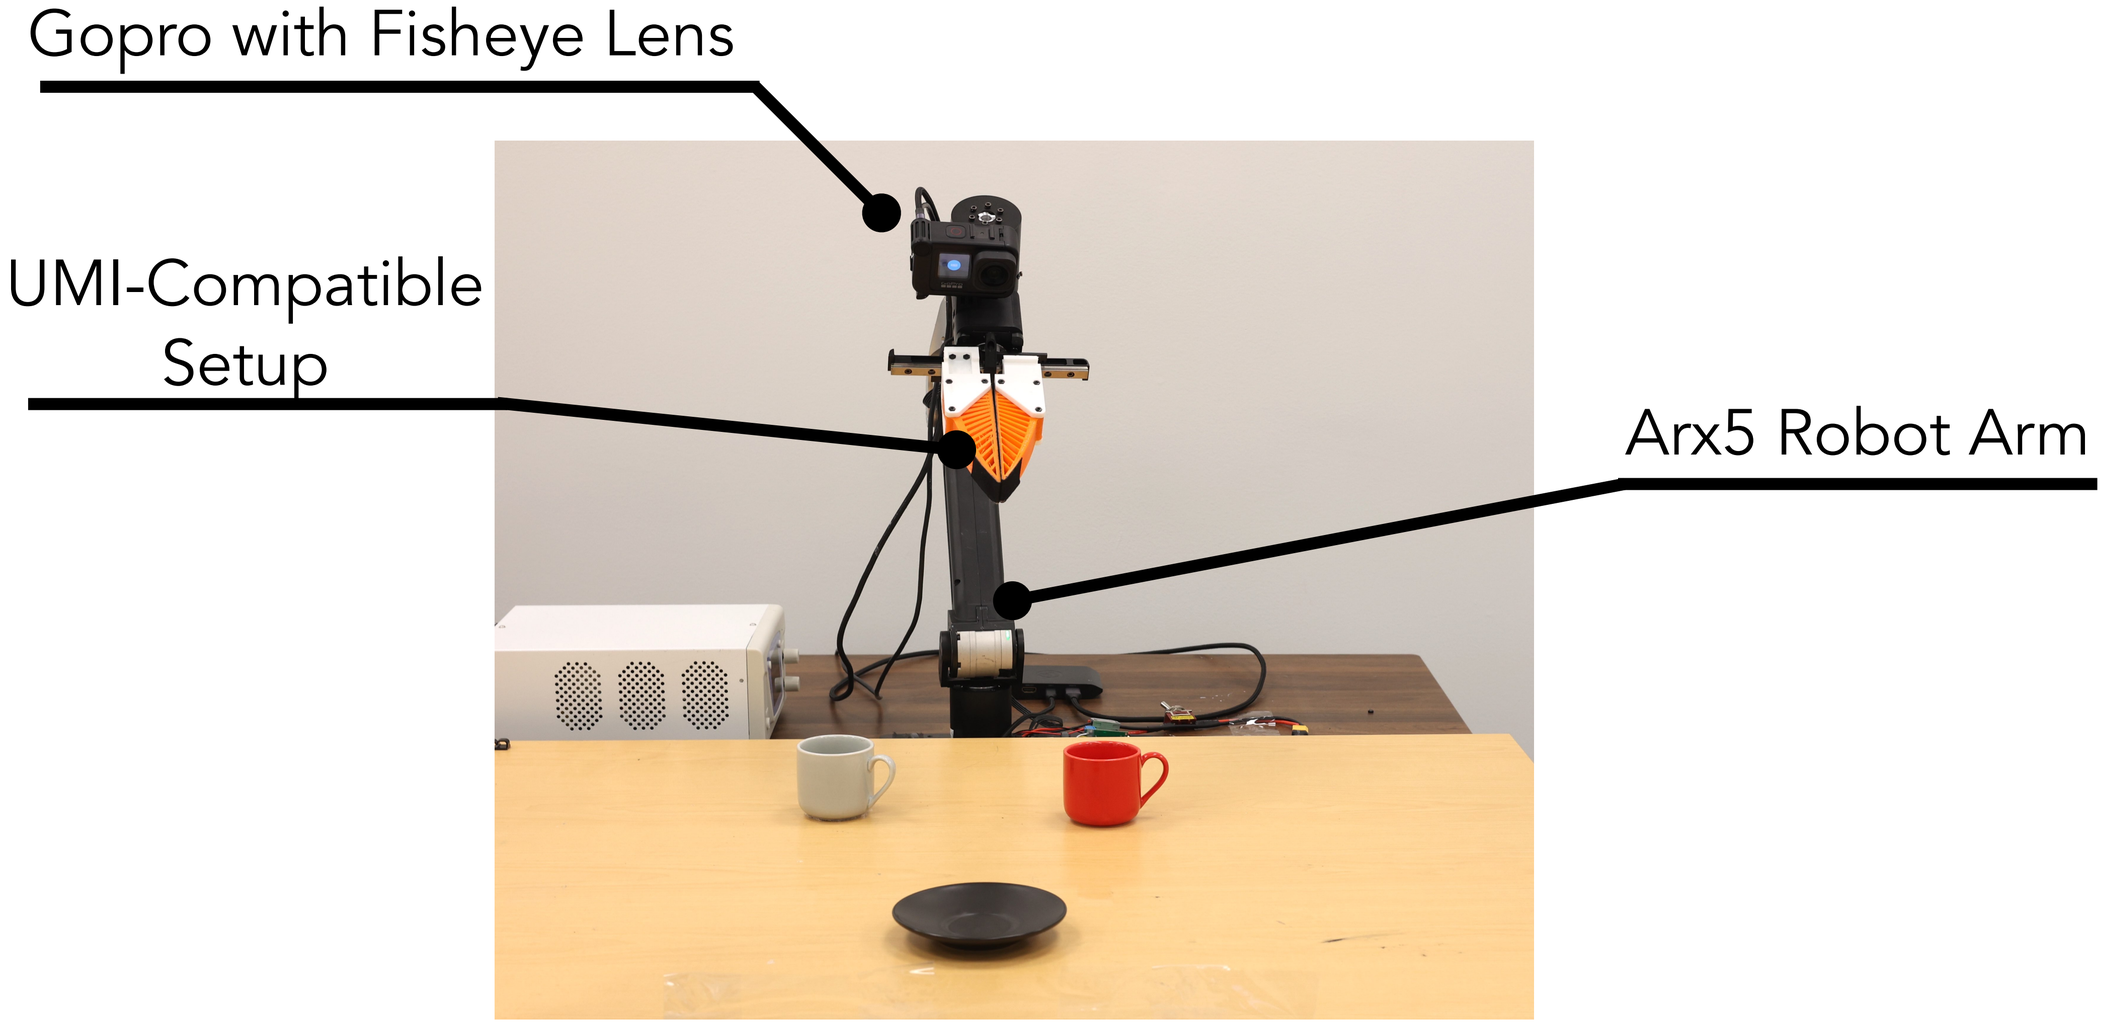

We use an off-the-shelf diffusion policy that arranges cups onto saucers, trained on the Universal Manipulation Interface (UMI) dataset & objects.

1) Can DynaGuide steer an off-the-shelf real robot policy?

Experimental Setup

Robot Perspective

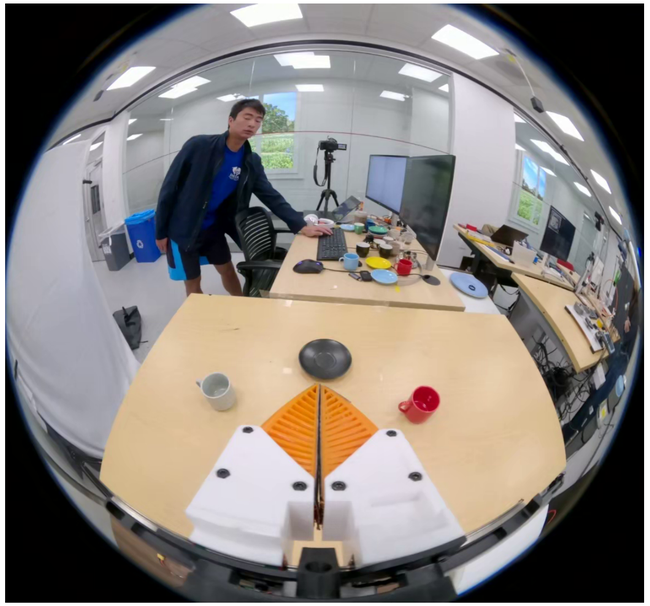

Guidance Condition Examples

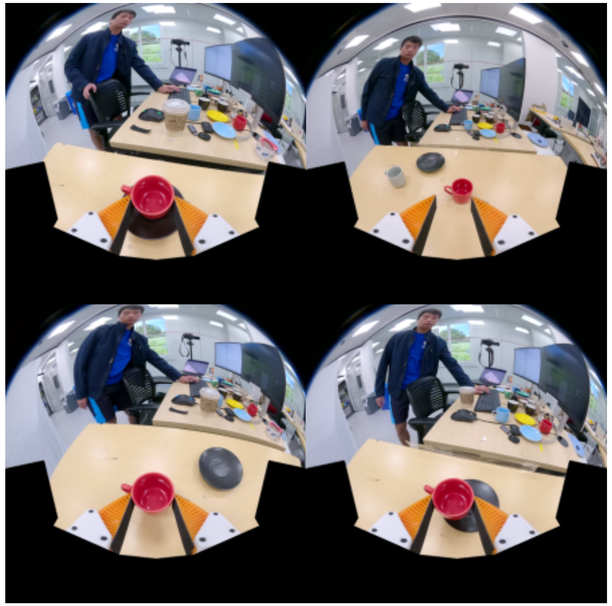

Investigate if DynaGuide can create cup color preference in the base policy

Paper Section: 4.5

1) Can DynaGuide steer an off-the-shelf real robot policy?

Base Policy

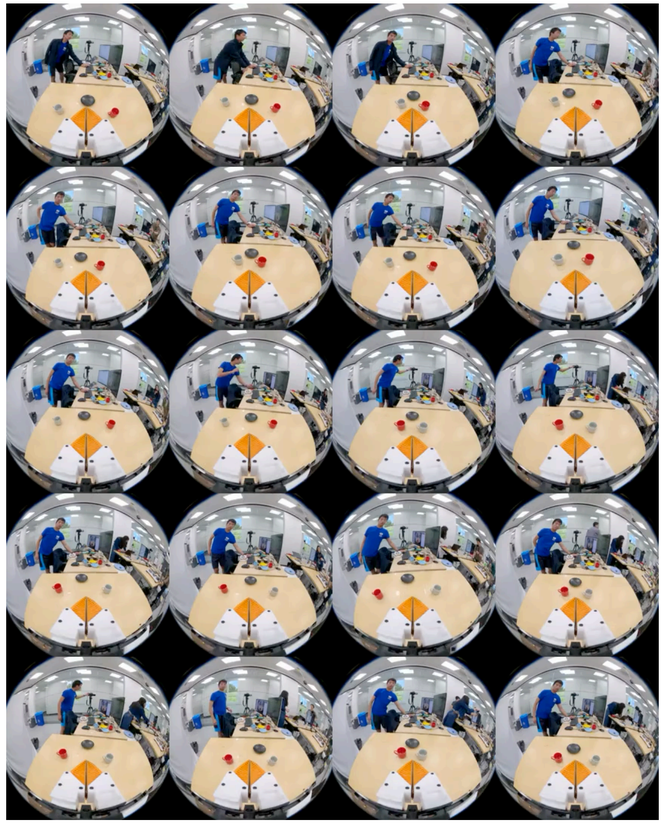

Red: 55%, Grey: 45%

Steer to Red

Red: 75%

Steer to Gray

Grey: 70%

DynaGuide successfully steers the base policy towards both colors of cups

Paper Section: 4.5

2) Can DynaGuide seek underrepresented behaviors on a real robot?

Experimental Setup

Robot Perspective

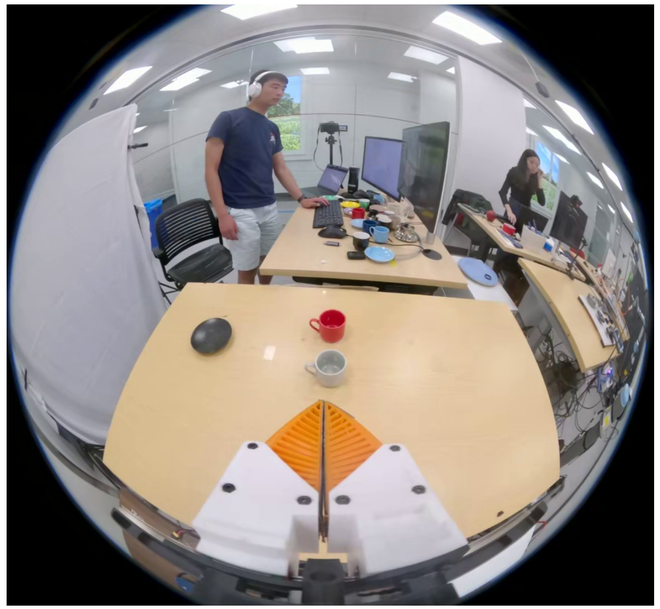

Guidance Condition Examples

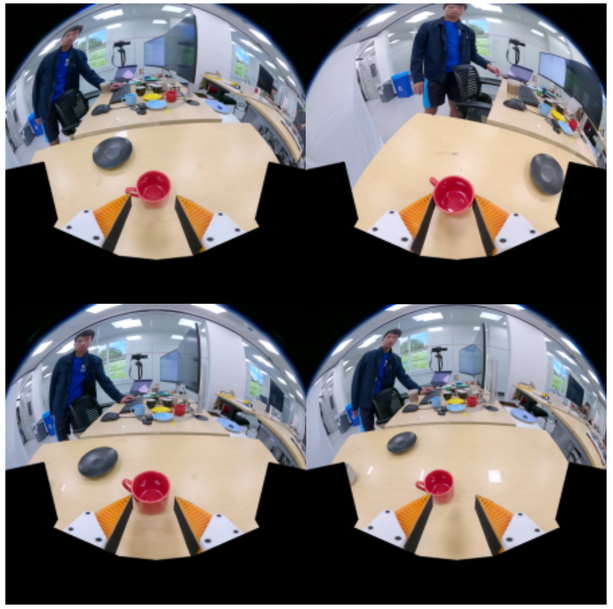

Investigate if DynaGuide can steer the base policy to avoid the closest cup.

Paper Section: 4.5

2) Can DynaGuide seek underrepresented behaviors on a real robot?

Base Policy

Red (far cup): 25%

Steer to Red

Red (far cup): 80%

Paper Section: 4.5

DynaGuide successfully steers the base policy towards an underrepresented behavior

Conclusion

DynaGuide…

Uses an **external dynamics model** to steer the behavior of a trained policy

Successfully extracts **single behaviors** from diverse policies and remains **robust** to lower quality objectives

Can **avoid** undesirable behaviors, steer towards **multiple** behaviors, and enhance **underrepresented** behaviors

Works on a **real robot** setup and successfully steers an **off-the-shelf** diffusion policy

DynaGuide: Steering Diffusion Polices with Active Dynamic Guidance

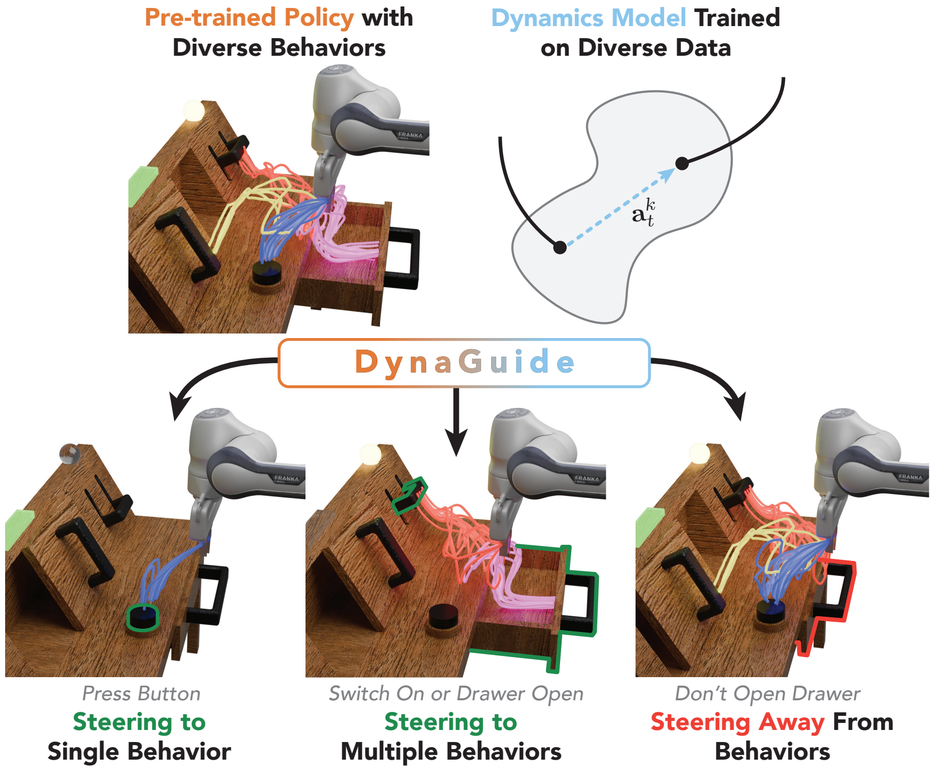

Maximilian Du, Shuran Song

Stanford Artificial Intelligence